LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

UCRL-CONF-208667

Verification and Validation of MERCURY: A Modern, Monte Carlo Particle Transport Code

R. J. Procassini, D. E. Cullen, G. M. Greenman, C. A. Hagmann

December 17, 2004

Monte Carlo 2005
Chattanooga, TN, United States
April 17, 2005 through April 21, 2005

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

The Monte Carlo Method: Versatility Unbounded In A Dynamic Computing World
Chattanooga, Tennessee, April 17–21, 2005, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2005)

VERIFICATION AND VALIDATION OF MERCURY: A MODERN, MONTE CARLO PARTICLE TRANSPORT CODE

Richard Procassini, Dermott Cullen,
Gregory Greenman and Christian Hagmann
Lawrence Livermore National Laboratory
Mail Stop L-95, P. O. Box 808
Livermore, CA 94551
spike@llnl.gov; cullen1@llnl.gov;
greenman@llnl.gov; hagmann1@llnl.gov

ABSTRACT

Verification and Validation (V&V) is a critical phase in the development cycle of any scientific code. The aim of the V&V process is to determine whether or not the code fulfills and complies with the requirements that were defined prior to the start of the development process. While code V&V can take many forms, this paper concentrates on validation of the results obtained from a modern code against those produced by a validated, legacy code. In particular, the neutron transport capabilities of the modern Monte Carlo code MERCURY are validated against those in the legacy Monte Carlo code TART.

The results from each code are compared for a series of basic transport and criticality calculations which are designed to check a variety of code modules. These include the definition of the problem geometry, particle tracking, collisional kinematics, sampling of secondary particle distributions, and nuclear data. The metrics that form the basis for comparison of the codes include both integral quantities and particle spectra. The use of integral results, such as eigenvalues obtained from criticality calculations, is shown to be necessary, but not sufficient, for a comprehensive validation of the code. This process has uncovered problems in both the transport code and the nuclear data processing codes which have since been rectified.

Key Words: Monte Carlo, particle transport, verification, validation

1 INTRODUCTION

Verification and Validation (V&V) is a critical, yet often overlooked, phase in the development cycle of any scientific computer code. These terms are similar, yet subtly different. *Verification* is "The process of determining whether or not the products of a given phase in the software life-cycle fulfill a set of established requirements" [1]. This implies an on-going process of unit testing to ensure that the *algorithms* which are implemented in the code solve the correct *equations* in order to calculate the required quantities. In contrast, *Validation* is "The stage in the software life-cycle at the end of the development process where software is evaluated to ensure that it complies with the requirements" [1]. This is a more comprehensive effort which is intended to test the *code in aggregate* to ensure that the code is obtaining the correct *results* for the required quantities.

This paper covers one aspect of the V&V process which is applied to test the neutron transport capabilities within the modern Monte Carlo code MERCURY [2]. These features in

MERCURY are validated through code-to-code comparisons with the legacy Monte Carlo code TART [3]. A series of basic transport and criticality calculations are run with each code in order to test various code modules, from problem geometry definition to particle tracking to collisional kinematics to the nuclear data used by the codes. The set of metrics that form the basis of comparison for the results from the two codes include both integral quantities (such as energy deposition in source problems or eigenvalues in criticality problems) and particle spectra (including the production, absorption and leakage spectra). This process of code-to-code comparison clearly shows that integral quantities are necessary, but not at all sufficient, for a comprehensive validation of a new code.

The organization of this paper is as follows. The features and capabilities of the MERCURY code are discussed in Section 2, while those of the TART code are presented in Section 3. A brief description of the overall V&V plan for the MERCURY code is given in Section 4. Section 5 presents the methodology of, and the results from, the code-to-code comparisons of MERCURY and TART. Finally, the summary and suggestions for future V&V studies are discussed in Section 6.

2 DESCRIPTION OF THE MERCURY MONTE CARLO TRANSPORT CODE

MERCURY is a modern, Monte Carlo particle transport code which has been developed at the Lawrence Livermore National Laboratory (LLNL) over the last six years. Funding to develop the code has come from the Advanced Simulation and Computing (ASC) and Science and Technology (S&T) programs at LLNL. MERCURY is envisioned as the eventual replacement for the legacy codes COG [4] and TART [3] as the next-generation Monte Carlo code at LLNL. It is our intent to maintain a multi-directorate code team into the future which will develop and support the code for use by the myriad of projects requiring particle transport simulations at LLNL.

The requirement of the ASC program to develop codes that can run on a variety of large-scale, parallel computing platforms has led to a three-pronged programming model in MERCURY. The three forms of parallelism supported in MERCURY are:

- *Domain Decomposition,* in which the problem geometry or mesh is spatially partitioned in order to support geometries with a large number of zones. This form of spatial parallelism is implemented via message passing methods.
- *Domain Replication,* in which the particle load is distributed across redundant copies of the spatial domain in order to support large numbers of particles. This form of particle parallelism is also implemented via message passing methods.
- *Task Decomposition*, in which the main particle loop is decomposed by assigning tasks (particle histories) to threads. This form of particle parallelism is implemented via shared-memory threading methods.

While MERCURY is written primarily in C, XML is used to describe the input data parameters during the parsing phase of a calculation. In the near future, the code will transition to use a small subset of the object-oriented features of the C++ programming language.

The current physics capabilities of MERCURY include:

- Time dependent transport of several types of particles through a background medium/ geometry, including (a) neutrons (n), (b) gammas (γ), and (c) the five lightest charged ions $({}^1H, {}^2H, {}^3H, {}^3He, {}^4He)$.
- Particle tracking in a wide variety of problem geometries, including (a) 1-D spherical meshes, (b) 2-D *r-z* meshes, (c) 3-D Cartesian meshes, (d) 3-D unstructured meshes, and (e) 3-D combinatorial geometry.
- Support for both multi-group and continuous energy treatment of cross sections.
- Population control can be applied to all types of particles. This capability is crucial for performing criticality calculations of subcritical or supercritical systems.
- Static k_{eff} and α eigenvalue, and pseudo-dynamic α eigenvalue "settle" calculations for criticality problems.
- Dynamic α (logarithmic population growth rate) calculations can be performed for any type of particles.
- Post-processing tally and diagnostic capabilities are provided by an auxiliary code named Caloris.
- Support for sources is rather limited at this time, but planned enhancements for these capabilities will made in the near future. The current capabilities include (a) external mono-energetic or fission spectrum sources, (b) external file-based sources and (c) zonal-based reaction sources.

In the near future, the following set of physics capabilities will be added to MERCURY:

- Each of the 7 types of particles will be able to interact with the background medium via (a) deposition of momentum, (b) deposition of energy and (c) depletion and production of isotopes resulting from nuclear reactions.
- The initial variance reduction capability to be added will be importance sampling, both with and without weight windows.
- The tally capabilities of MERCURY will be rewritten and generalized. In time, the code will also support event history tallies.

3 DESCRIPTION OF THE TART MONTE CARLO TRANSPORT CODE

TART [3] is a Monte Carlo particle transport code which has a long history of use at LLNL. The development of TART was begun in the early 1960s by Ernest Plechaty. Some of its current applications at LLNL include NIF shielding calculations and criticality safety.

The physics capabilities of the TART code include:

- Time dependent transport of (a) neutrons (n) and (b) gammas (γ) through a background medium/geometry.
- Particles are tracked though 3-D combinatorial geometries.

- Only a multi-group treatment of cross sections is supported. However, the code also includes (a) a multi-band statistical treatment of resolved resonances and (b) a separate statistical treatment of unresolved resonances.
- Population control, for use in criticality problems, can be applied to both neutrons and gammas.
- Static k_{eff} and α eigenvalue, and pseudo-dynamic α eigenvalue "settle" calculations for criticality problems.
- A wide variety of flexible particle sources and tallies are supported.

4 OVERVIEW OF THE MERCURY VERIFICATION AND VALIDATION PLAN

The plan for the Verification and Validation (V&V) of the MERCURY code is divided into three main areas. The first area of V&V focuses with the calculation of benchmark test problems. The set of problems that are target for calculation by MERCURY include:

- Analytic problems compiled by Sood, Forster and Parsons [5], Kobayashi, Sugimura and Nagaya [6], etc.
- Experimental criticality problems compiled in the *International Handbook of Evaluated Criticality Safety Benchmark Experiments* (ICSBEP Handbook) [7].
- Time-dependent transport problems performed as part of the LLNL pulsed spheres experimental program [8].

The second area of V&V involves code-to-code comparisons against other Monte Carlo codes which have been previously validated. The codes which are currently planned for comparison with MERCURY include TART [3], COG [4] and MCNP [9]. This effort will compare:

- Integral quantities, such as k_{eff} and α eigenvalues in criticality calculations and tallies of energy deposition or dose rate in source problems.
- Particle spectra such as production, absorption and leakage

The third area of V&V involves numerical resolution studies on a subset of the problems listed above. This is intended to *verify* the convergence of the the code's results as the spatial, temporal and energy resolution of the problem is increased.

The balance of this paper will focus on code-to-code comparisons between MERCURY and TART for a series of basic transport ans criticality problems.

5 CODE-TO-CODE COMPARISON OF MERCURY AND TART RESULTS

The transport capabilities of TART have been previously validated against many of the benchmark test problems listed in the previous section, as well as through comparisons with several other Monte Carlo codes [10], including MCNP [9], KENO [11], VIM [12] and COG [4]. We rely on previous validation of TART in order to V&V the neutron transport capabilities of MERCURY through code-to-code comparisons.

Accurate code-to-code comparisons of MERCURY and TART require minimizing the differences in the two models. Hence, for this study each of the codes is run in the following manner.

In an effort to ensure that differences do not arise from particle tracking, combinatorial geometry is used to model the system of interest. In order to minimize any differences to to statistics, the same number of particle histories will be run by each code.

The same set of nuclear data is used in each code. The point wise data that is converted into the multi-group constants used by the code is a hybrid of two evaluated data sets, where (a) cross sections (σ) are obtained from the ENDF/B-VI (Release 8) evaluation and (b) secondary particle distributions $(E \rightarrow E', \Omega \rightarrow \Omega')$ are obtained from the from ENDL-94. While the underlying point wise nuclear data is the same for these calculations, it should be pointed out that different processing codes were used to generate the binary files that contain group constants. The PREPRO [13] package is used to generate TART's data files, while MCFGEN serves the same purpose for MCAPM [14], the cross section server and collisional kinematics package that is used by MERCURY.

The same energy treatment of the nuclear data is used by each of the codes. All of the problems presented here employ a multi-group treatment of the cross sections with a 616-groups. The group boundaries are chosen to be equally spaced in lethargy $u = \ln(E_{max}/E)$, with 50 groups per decade over the range $1.0\times10^{-11} < E_n < 20$ MeV.

5.1 Basic Transport Calculations

The rudimentary transport capabilities of *any* transport code can be validated through code intercomparison of a "Broomstick". The Broomstick is a long, thin rod composed of a single isotope material. Particles from a parallel-ray, monoenergetic source are injected down the axis of the rod. The length of the rod is chosen to ensure that *all* particles collide before reaching the far end, while the radius of the rod is chosen to ensure that particles only collide *once* before leaking from the system. This geometry is chosen such that the particle leakage spectrum becomes the principal diagnostic, since it directly corresponds to the secondary energy spectrum of the collided particles.

The Broomstick problem may be used to validate many portions of a Monte Carlo transport code, including (a) various types of particle sources, such as disk, point, spherical volume, etc., (b) sampling of particle mean-free paths and the distance to collision, and (c) collisional kinematics and the sampling of secondary particles distributions within the collision package.

The Broomstick used in this study, shown in Figure 1, is defined as follows:

- The dimensions of the rod are length $L = 10^5$ cm and radius $r = 10^{-5}$ cm.
- The rod is composed of pure ^{235}U, at a density of $\rho = 18.7$ g/cm^3 and temperature of $T = 300$ K.
- A monoenergetic, parallel-ray disk source of radius $r_{source} = 0$ is directed down the axis of the rod.

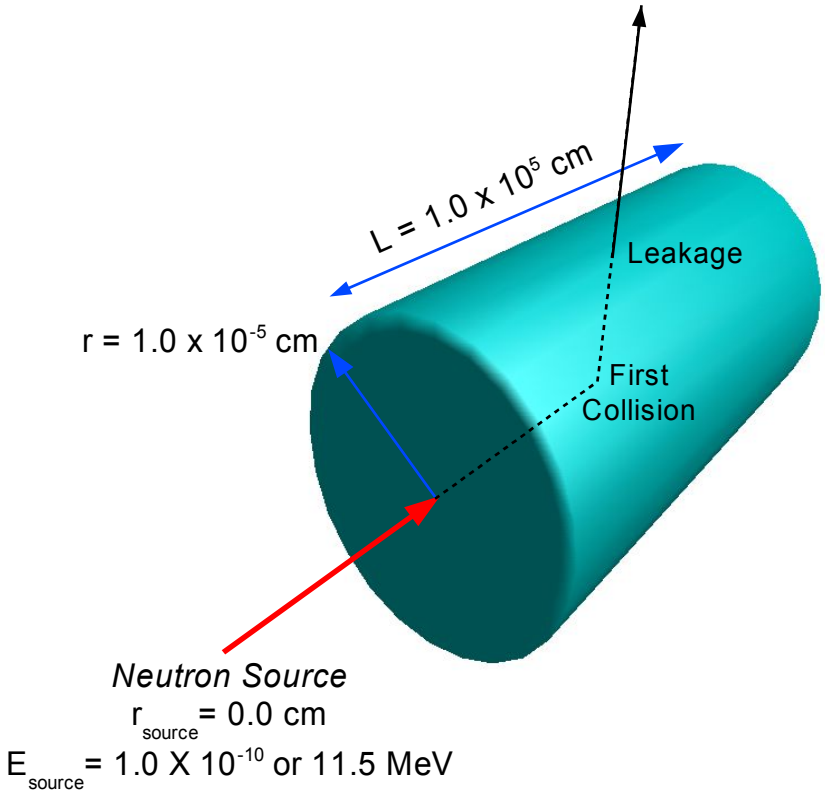

Figure 1. The geometry of the Broomstick problem.

Two variants of this Broomstick problem was studied in which only the source energy was varied. The intent is to investigate different collisional interactions with each problem:

- The low energy ($E_{source} = 1.0 \times 10^{-10}$ MeV) case has only two reaction channels (a) elastic scattering (n, n), and (b) fission (n, f).
- The high energy ($E_{source} = 11.5$ MeV) case has three reaction channels (a) elastic (n, n) and inelastic (n, n') scattering, (b) multiple particle production $(n, 2n)$, and (c) fission (n, f).

The number of particles injected into the rod in each variant is $N_{sim} = 5 \times 10^8$.

5.1.1Low-Energy Broomstick Problem

The particle leakage spectrum from the low-energy Broomstick problem is shown in Figure 2. The MERCURY results are represented by the blue curve, while the TART results are shown as the red curve. The black curve is the point wise, secondary-particle fission energy spectrum from the ENDL-94 evaluation. This data is tabulated at several energies, including $E_n = 1.0 \times 10^{-10}$ and $E_n = 11.5$ MeV.

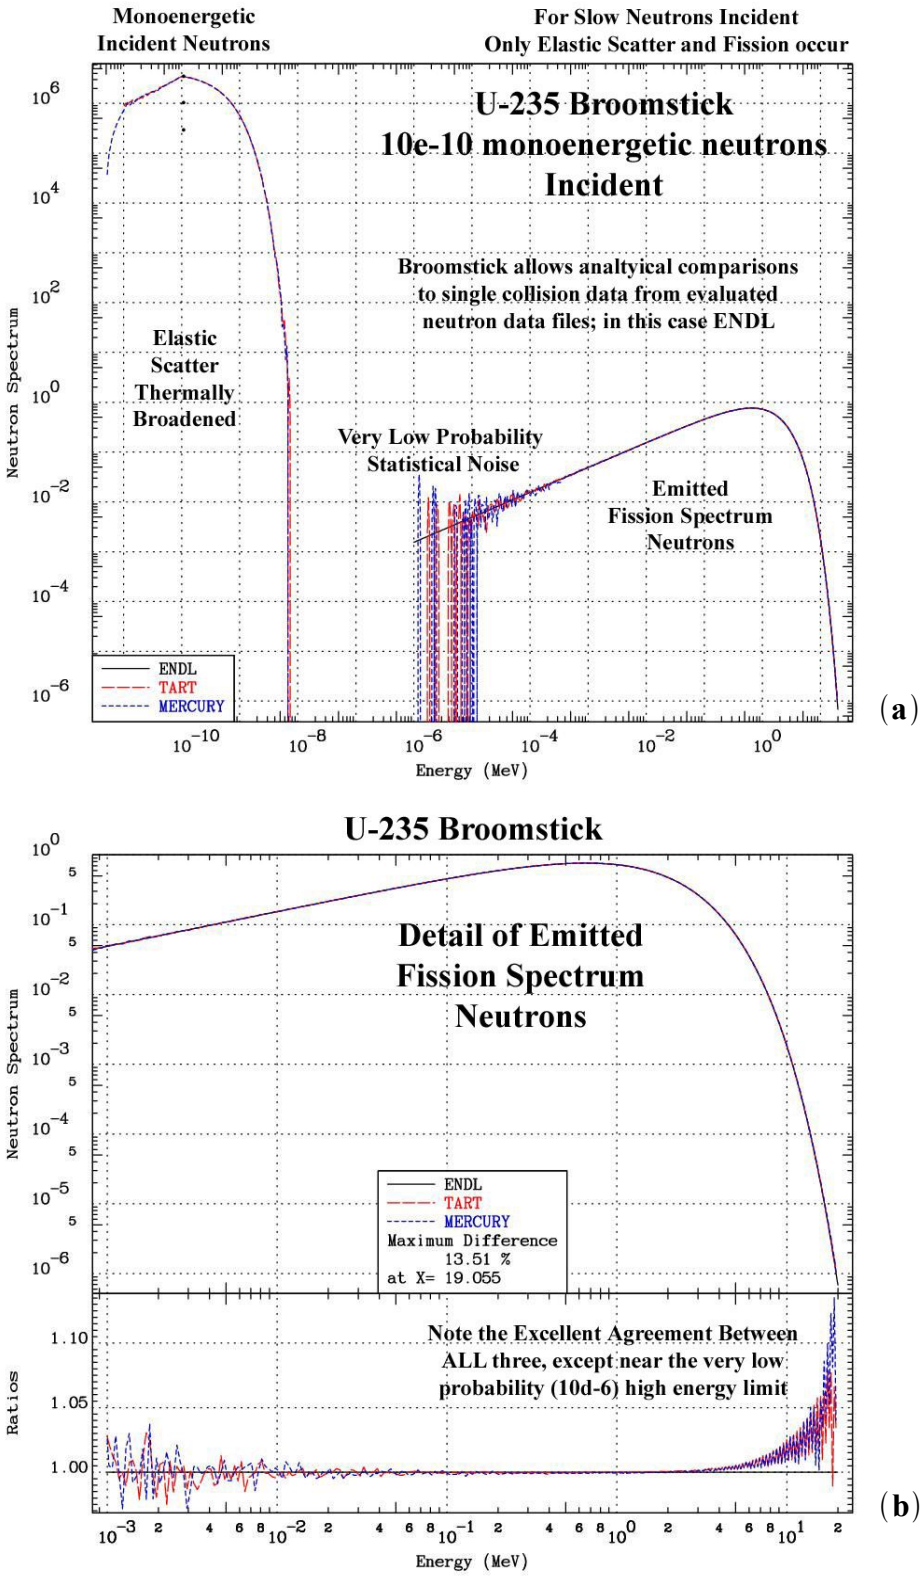

Figure 2. The particle leakage spectrum from the low-energy ($E_{source} = 1.0\times10^{-10}$ MeV) Broomstick problem: (a) full energy range, (b) high energy range. The ENDL-94 data is shown in black, MERCURY results are blue and TART results are red.

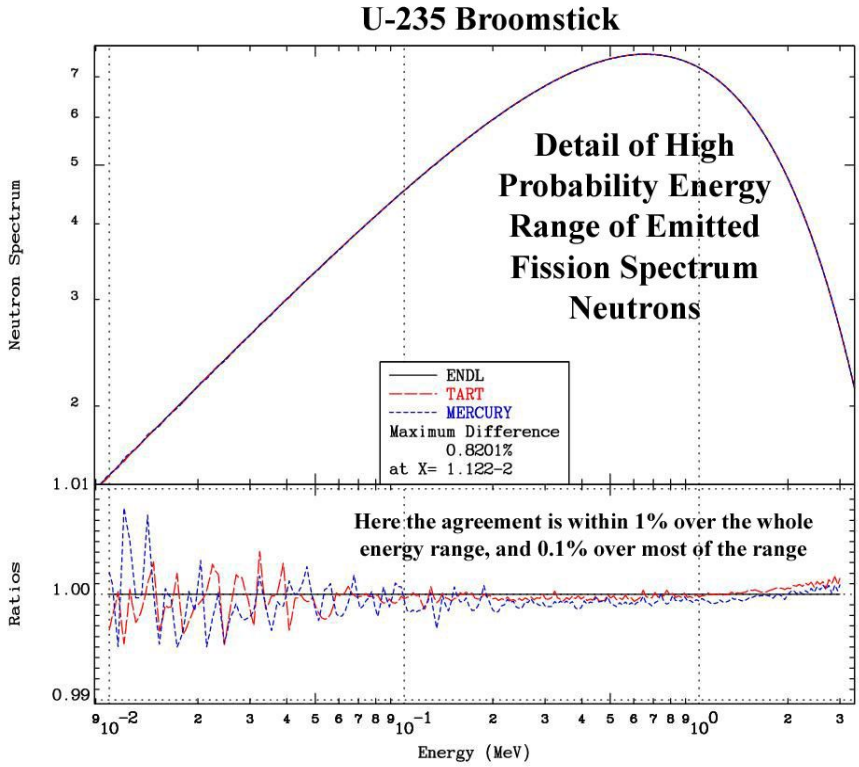

(c)

Figure 2 (continued). The particle leakage spectrum from the low-energy ($E_{source} = 1.0\times10^{-10}$ MeV) Broomstick problem: (c) high-energy, high-probability range. The ENDL-94 data is shown in black, MERCURY results are blue and TART results are red.

Notice how the neutron leakage spectrum shown in Figure 2a is divided into two regions. The low energy ($1.0\times10^{-11} < E_n \le 7.0\times10^{-9}$ MeV) regions results from elastic scattering off of the ^{235}U nuclei at a temperature of $T = 300$ K ($T = 2.53\times10^{-8}$ MeV). The incident energy at which the neutrons were injected into the rod is shown by the dots at $E_n = 1.0\times10^{-10}$ MeV. Figure 2a clearly shows excellent agreement of the particle leakage spectra from both MERCURY and TART with the ENDL-94 fission spectrum. The main differences occur in the low-probability range of the fission spectrum for $1.0\times10^{-6} < E_n < 3.0\times10^{-4}$ MeV.

Figure 2b and 2c show that the MERCURY and TART results are within 2% of the ENDL-94 data, except at the high-energy range of the fission spectrum for $3 < E_n < 20$ MeV, where the differences approach 13%. Over the energy range of $1.0\times10^{-3} < E_n < 3$ MeV, the difference between the predicted results and the underlying nuclear data is less that 0.1%.

5.1.2High-Energy Broomstick Problem

The secondary-neutron energy spectrum for $E_{source} = 11.5$ incident energy neutrons is shown in Figure 3. This shows the secondary-energy distributions for the three possible reactions (excluding elastic scattering) in ^{235}U at the incident energy. This data is taken directly from the point wise ENDL-94 evaluation. Fission (n, f) is shown in red, $(n, 2n)$ is shown in green ans inelastic scattering (n, n') is shown in blue. The sum of these three curves is shown in black.

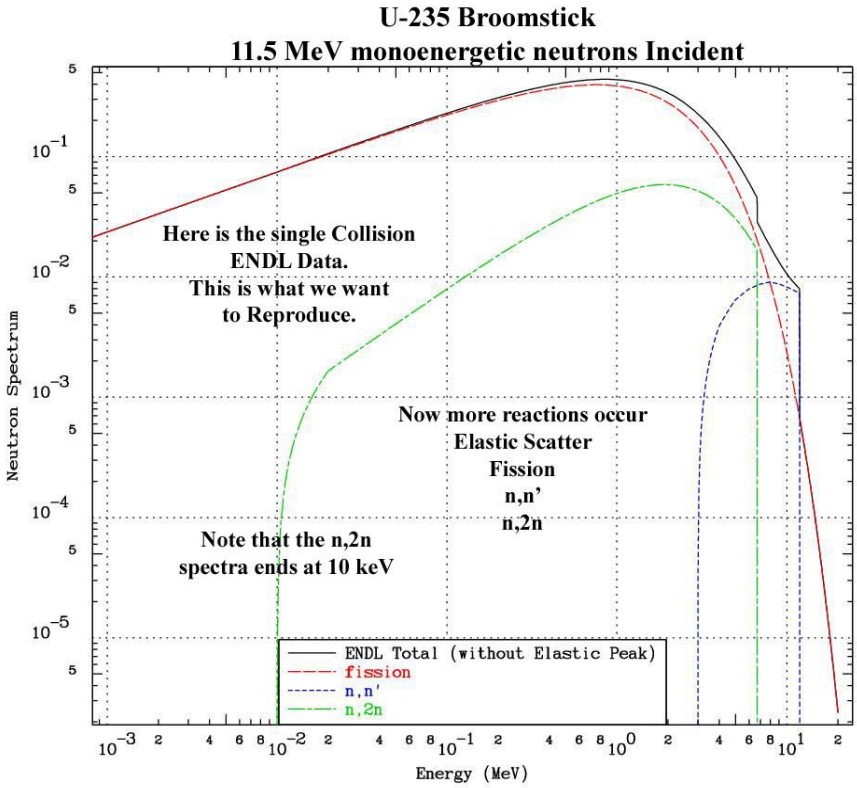

Figure 3. The secondary-energy spectra for the low-energy (E_{source} = 11.5 MeV) Broomstick problem, taken from the ENDL-94 evaluation, showing the fission (n,f) spectrum (red), the $(n,2n)$ spectrum (green), the inelastic scattering (n,n') spectrum (blue), and the sum (black).

The particle leakage spectrum from the high-energy Broomstick problem is presented in Figure 4. The same color schemes that was used in Figure 2 is also used here. The ENDL-94 summed-spectra curve from Figure 3 is shown in black. The agreement between TART and the ENDL-94 data is excellent, except for the low-energy range $E_n < 1.0\times10^{-3}$ MeV, which is dominated by statistical noise, and the elastic scattering peak centered at $E_n < 11.5$, where the effect of elastic scattering is not included in the black curve.

The agreement between the MERCURY and TART results is very good, with the exception of a bump in the MERCURY leakage spectrum in the energy range $10 < E_n < 100$ keV (see Figures 4a and 4b). This difference arises from the method of sampling particle energies from the secondary particle distribution, in this case for $(n,2n)$. TART samples the secondary particle energy from equally-probable (histogrammed) bins, with the exception of the highest and lowest energy bins, where linear interpolation was used. At the time these results were obtained, MERCURY used only equally-probable bin sampling for the secondary particle energies. The result is the observed bump in the spectrum, which occurs when the distribution falls off rapidly in the first or last bin, as is the case for a threshold reaction such as $(n,2n)$. This difference in sampling has since been fixed in the MCAPM library.

This results shows the power of using simple test cases, such as the monoenergetic Broomstick problem, for V&V purposes.

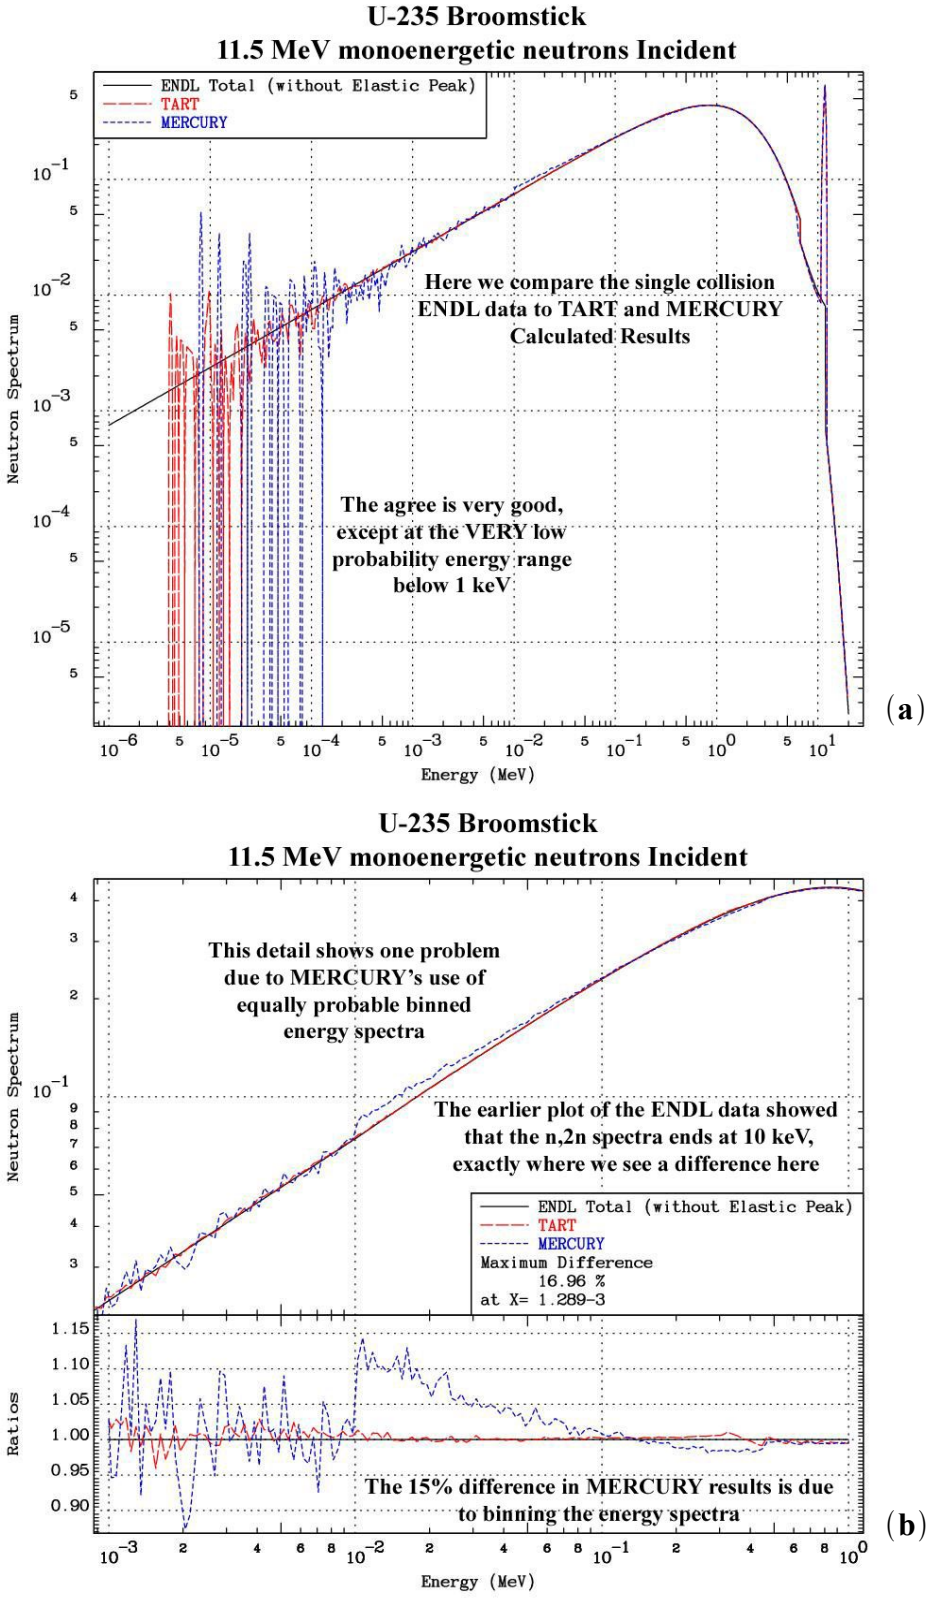

Figure 4. The particle leakage spectrum from the high-energy (E_{source} = 11.5 MeV) Broomstick problem: (a) full energy range, (b) high-probability, intermediate-energy range. The ENDL-94 data is shown in black, MERCURY results are blue and TART results are red.

5.2 Criticality Calculations

For this portion of the study, static k_{eff} eigenvalue calculations of two fast critical assemblies and two thermal systems are performed by each of the codes. In this method, all of the neutrons in the system at the start of a generation are tracked unto they are removed from the system via leakage or absorption. The secondary particles that result from reactions with the absorbed neutrons form the source for the next generation. The method iterates over multiple generations, in both a transient and equilibrium phase, in order to produced a "settled" particle distribution, and estimates of the k_{eff} and α eigenvalue for the system.

Code-to-code comparisons are made for both integral results, as well as particle spectra. The integral parameters include the k_{eff} and α eigenvalues, and the neutron removal lifetime τ_{rem}. The eigenvalues are defined as:

$$k_{eff} = \frac{N_{prod}}{N_{abs} + N_{leak}} \tag{1}$$

$$\alpha = \left(\frac{1}{\tau_{rem}}\right)\left(\frac{N_{prod}}{N_{abs} + N_{leak}} - 1\right) = \left(\frac{1}{\tau_{rem}}\right)\left(k_{eff} - 1\right) \tag{2}$$

where N_{prod}, N_{abs} and N_{leak} are the number of particles produced, absorbed and leaked during the generation, and τ_{rem} is the generation-averaged removal lifetime. The particle spectra compared include the production, absorption and leakage spectra. Note that the absorption and production spectra represent the *incident* energy of particles that are absorbed and produce secondary particles, respectively.

5.2.1Fast Critical Assemblies

The two fast critical assemblies that are modeled in this study are taken from the ICSBEP Handbook [7]. One uranium and one plutonium system were modeled. The Godiva assembly, which is given the moniker HEU-MET-FAST-001, is a bare (unreflected), oralloy (highly enriched uranium) system. The Jezebel assembly, known as PU-MET-FAST-001, is a bare, δ-phase plutonium system. These calculations were performed with $N_{sim} = 1\times10^{5}$ particles per generation, and with a convergence criterion on the static k_{eff} eigenvalue calculation of $\epsilon = 1.0\times10^{-4}$.

Integral Parameter Results

The integral results for the Godiva and Jezebel critical assemblies are presented in Tables I through III. Table I shows that MERCURY and TART are each calculating the same k_{eff} eigenvalue to within the specified convergence criteria. This results is to be expected, since each of the calculations were run to convergence, where the quantity that is checked for convergence is the k_{eff} eigenvalue. The differences in the neutron removal lifetime τ_{rem} calculated by the two codes (see Table II) is about an order of magnitude larger than convergence criteria, but the results are still in very good agreement.

Table I. Code-to-code comparisons of the k_{eff} eigenvalue.

System	TART	MERCURY	Difference (TART – MERCURY)
Godiva	1.00492	1.00513	-0.02%
Jezebel	1.00138	1.00149	-0.01%

Table II. Code-to-code comparisons of the neutron removal lifetime τ_{rem}.

System	TART (μ sec)	MERCURY (μ sec)	Difference (TART – MERCURY)
Godiva	6.051×10^{-3}	6.061×10^{-3}	-0.16%
Jezebel	3.910×10^{-3}	3.923×10^{-3}	-0.33%

Table III. Code-to-code comparisons of the α eigenvalue.

System	TART (gen/ μ sec)	MERCURY (gen/ μ sec)	Difference (TART – MERCURY)
Godiva	8.128×10^{-1}	8.463×10^{-1}	-4.12%
Jezebel	3.530×10^{-1}	3.806×10^{-1}	-7.82%

The level of agreement between the α eigenvalue calculated by MERCURY and TART is not nearly as good as that for k_{eff} or τ_{rem}. However, this is not unexpected. When k_{eff} is close to unity, being calculated to a high precision, one should expect that α will have a large uncertainty, since $\alpha = \left(k_{eff} - 1\right)/\tau_{rem}$ and $\left(k_{eff} - 1\right) \ll 1$. In general, the accuracy of the calculated values of α improves as k_{eff} moves away from unity (in either direction).

Taking into account the large uncertainty in α when $k_{eff} \simeq 1$, these integral results seem to suggest that the MERCURY is in agreement with TART, and working correctly. Correct? Let us take a close look in the form of the particle spectra.

Particle Spectra

The production spectrum for the Godiva critical assembly is shown in Figure 5. Our initial MERCURY calculation of Godiva (shown in red) exhibits significant differences from the TART results (shown in black). MERCURY has much less production for incident neutrons with intermediate energies $E_n < 10$ keV and for high energies $E_n > 3$ MeV.

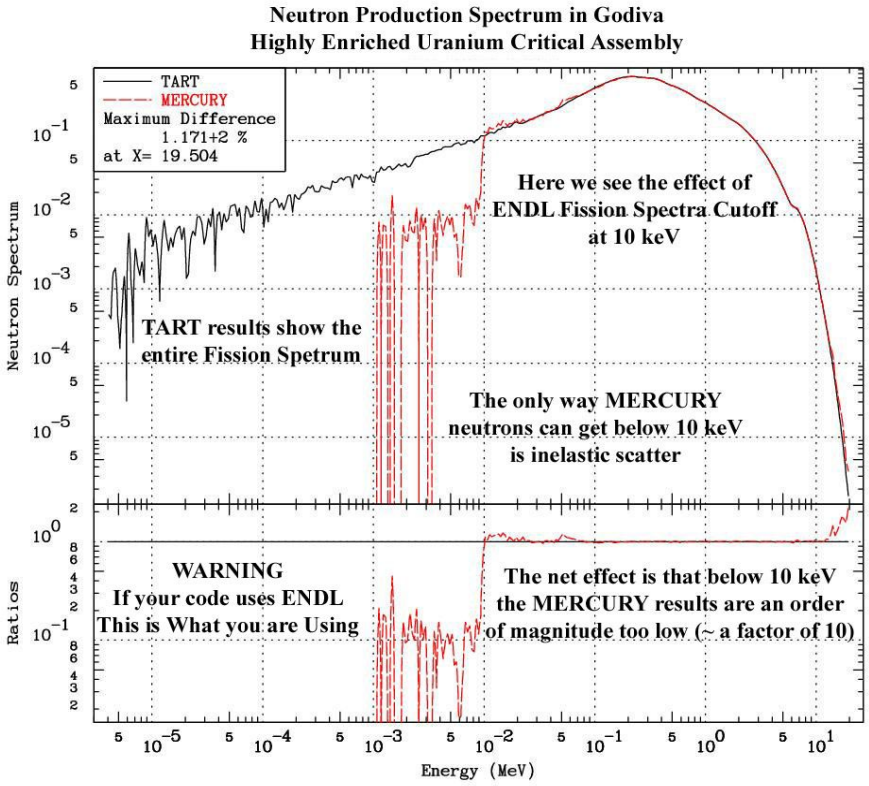

(**a**)

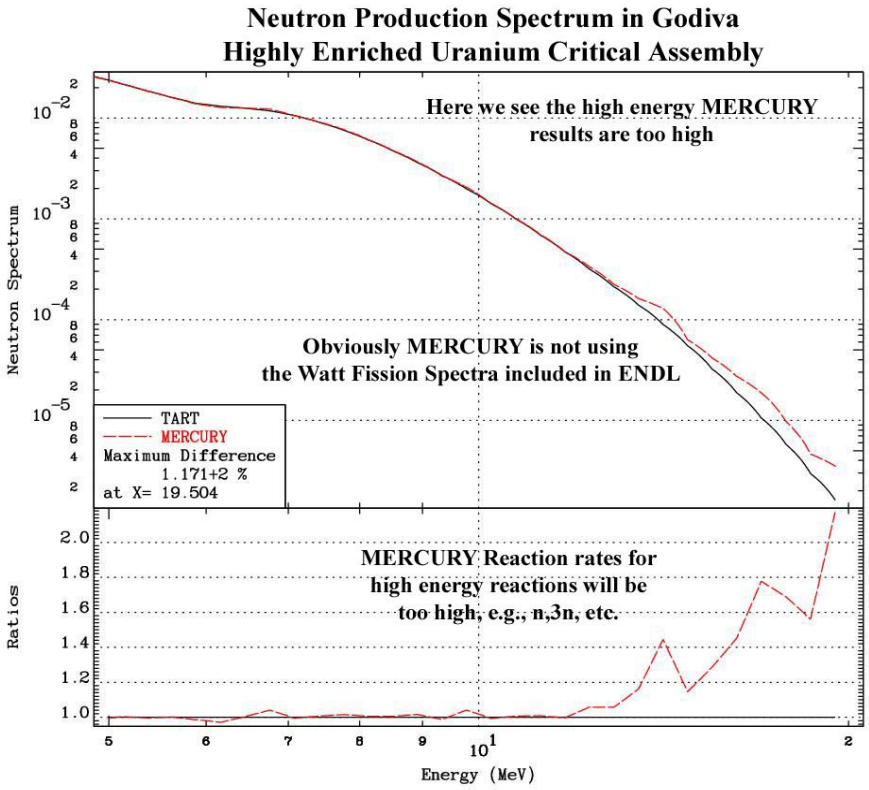

(**b**)

Figure 5. The particle production spectrum from the Godiva criticality problem. Significant differences in the spectra from MERCURY (red) and TART (black) (a) for energies $E_n < 10$ keV, (b) for energies $E_n > 3$ MeV.

Both of these effects were traced back to the method of sampling fission spectrum neutrons in the MCAPM package. While TART samples a target-mass-parameterized Watt spectrum to obtain secondary fission neutron energies (see the blue and black curved in Figure 6), MCAPM/ MERCURY sampled tabulated data which (a) terminated at $E_n = 10$, so no fission neutrons were created at energies below 10 keV, and (b) used an inaccurate equally-probable bin sampling algorithm at energies above a few MeV. While the tabulated data used by MCAPM agree with the Watt spectrum in the energy range $1.0 \times 10^{-2} < E_n < 3$ MeV, the low and high energy tails were not being sampled correctly.

Once MCAPM was modified to sample from the same Watt representation of the fission spectrum that TART uses, the agreement between the two codes for Godiva is excellent, as is shown in Figure 7. It should be pointed out that while the effect of changing the sampling of the fission spectrum resulted in large changes to the production spectrum, the changes to the integral results presented in Tables I through III were within the error bars shown in the Tables. This is an important lesson learned. While the comparison of integral results are necessary component of a V&V plan for a transport code, such comparisons are not sufficient to ensure that the code is working correctly. A detailed comparison of the particle production, absorption and leakage spectra is also required. Once all of the integral and spectral results of the two codes are in agreement, one can be fairly confident that the codes are working correctly.

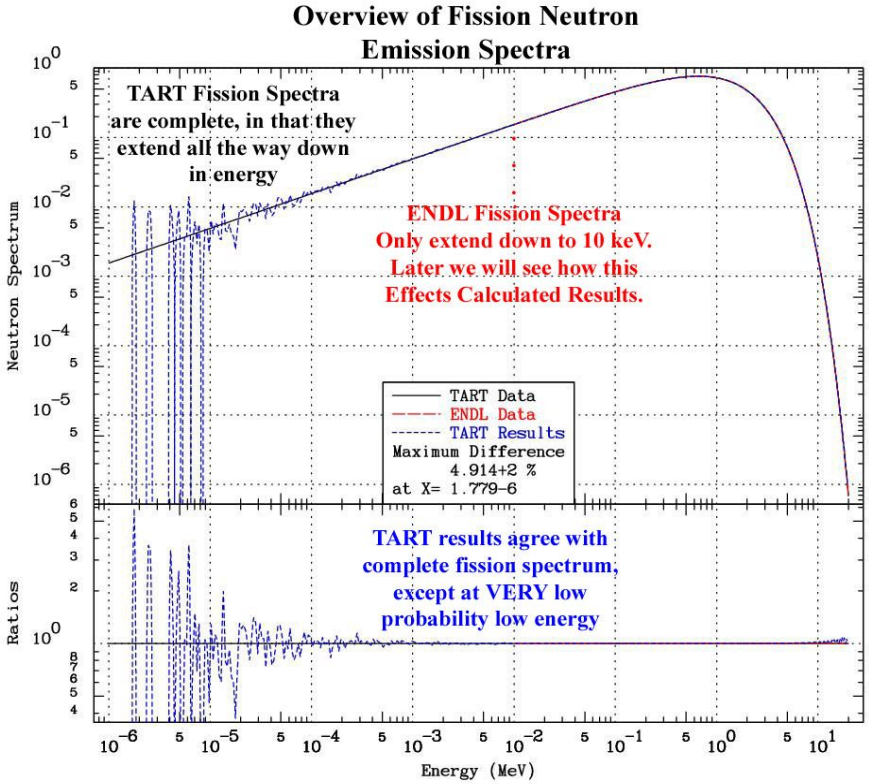

Figure 6. The secondary-energy fission neutron spectrum as tabulated in the ENDL-94 data evaluation (red), in the form of a Watt spectrum in the TART data files (black), and as sampled by TART (blue). The red dots at 10 keV indicate where the ENDL-94 data terminates.

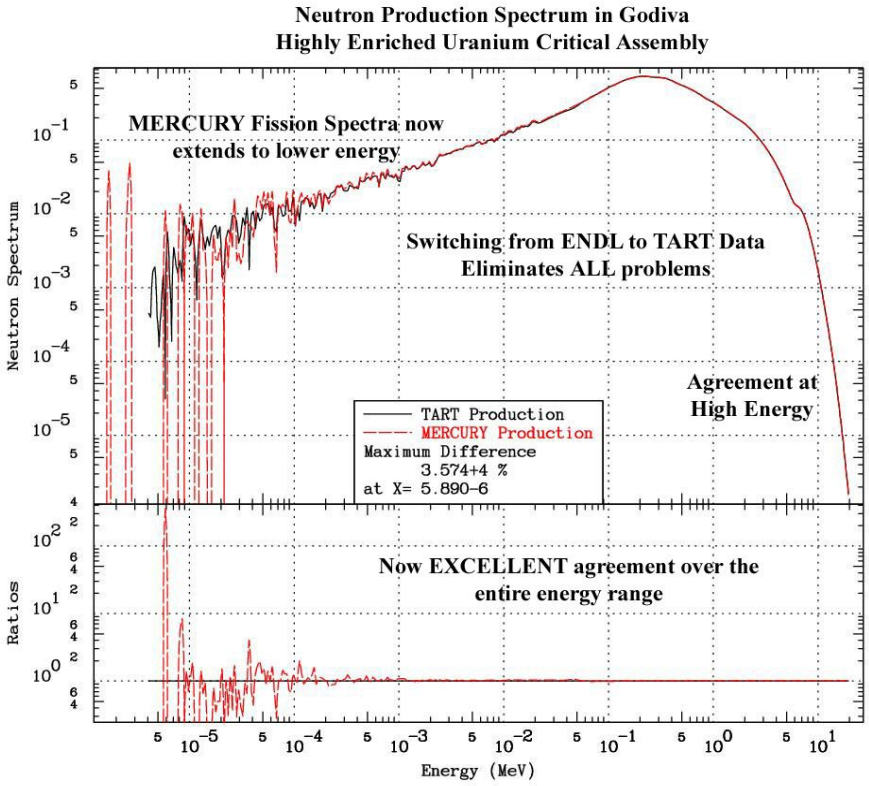

Figure 7. The particle production spectrum from the Godiva criticality problem. The modified version of MCAPM which uses a Watt fission spectrum yields excellent agreement between MERCURY (red) and TART (black).

The production spectrum for the Jezebel critical assembly is shown in Figure 8. Our initial MERCURY calculation of Jezebel (the red curve in Figure 8a) is markedly different from the TART results (shown in black). MERCURY has significantly more production for incident neutrons with intermediate energies $E_n < 150$ keV.

This effect has also been traced to the sampling of secondary particle energies from equally-probable bins in MERCURY, which produced the bump in the $(n,2n)$ distribution shown in Figure 4. In the case of ^{239}Pu, the lowest-energy equally probable bin in the ENDL-94 evaluation extends from $E_n = 0$ to $E_n = 150$ keV, and the old MCAPM implementation samples particle energies uniformly over that entire energy interval. Once the sampling method was modified to use linear-interpolation in the lowest and highest energy bins, the excellent agreement shown in Figure 8b was obtained.

5.2.2Thermal Systems

Particle Spectra

The Godiva and Jezebel critical assemblies were modified to create two thermal systems by mixing 1 part of fissile material in 100 parts of water. These homogeneous solutions, known as "Wet Godiva" and "Wet Jezebel", are important tests of the elastic-scattering thermalization model that has recently been implemented in MCAPM/MERCURY.

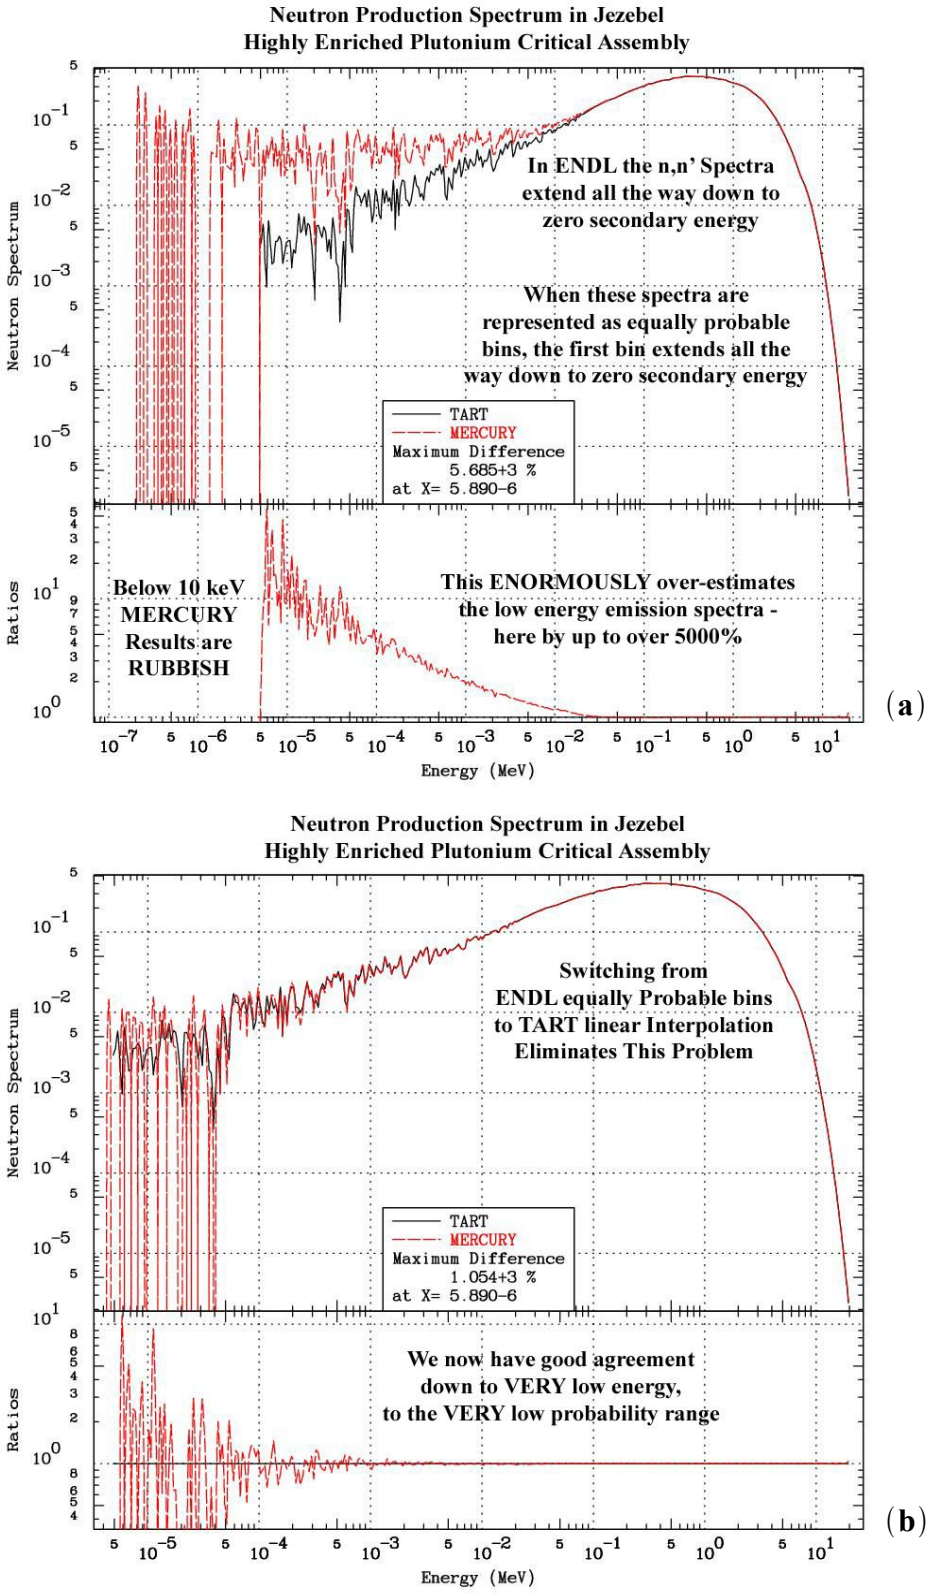

Figure 8. The particle production spectrum from the Jezebel criticality problem: (a) Significant differences in the spectra from MERCURY (red) and TART (black) are observed for energies $E_n < 150$ keV, (b) Excellent agreement is obtained when linear interpolation of secondary particle energies is used in MERCURY.

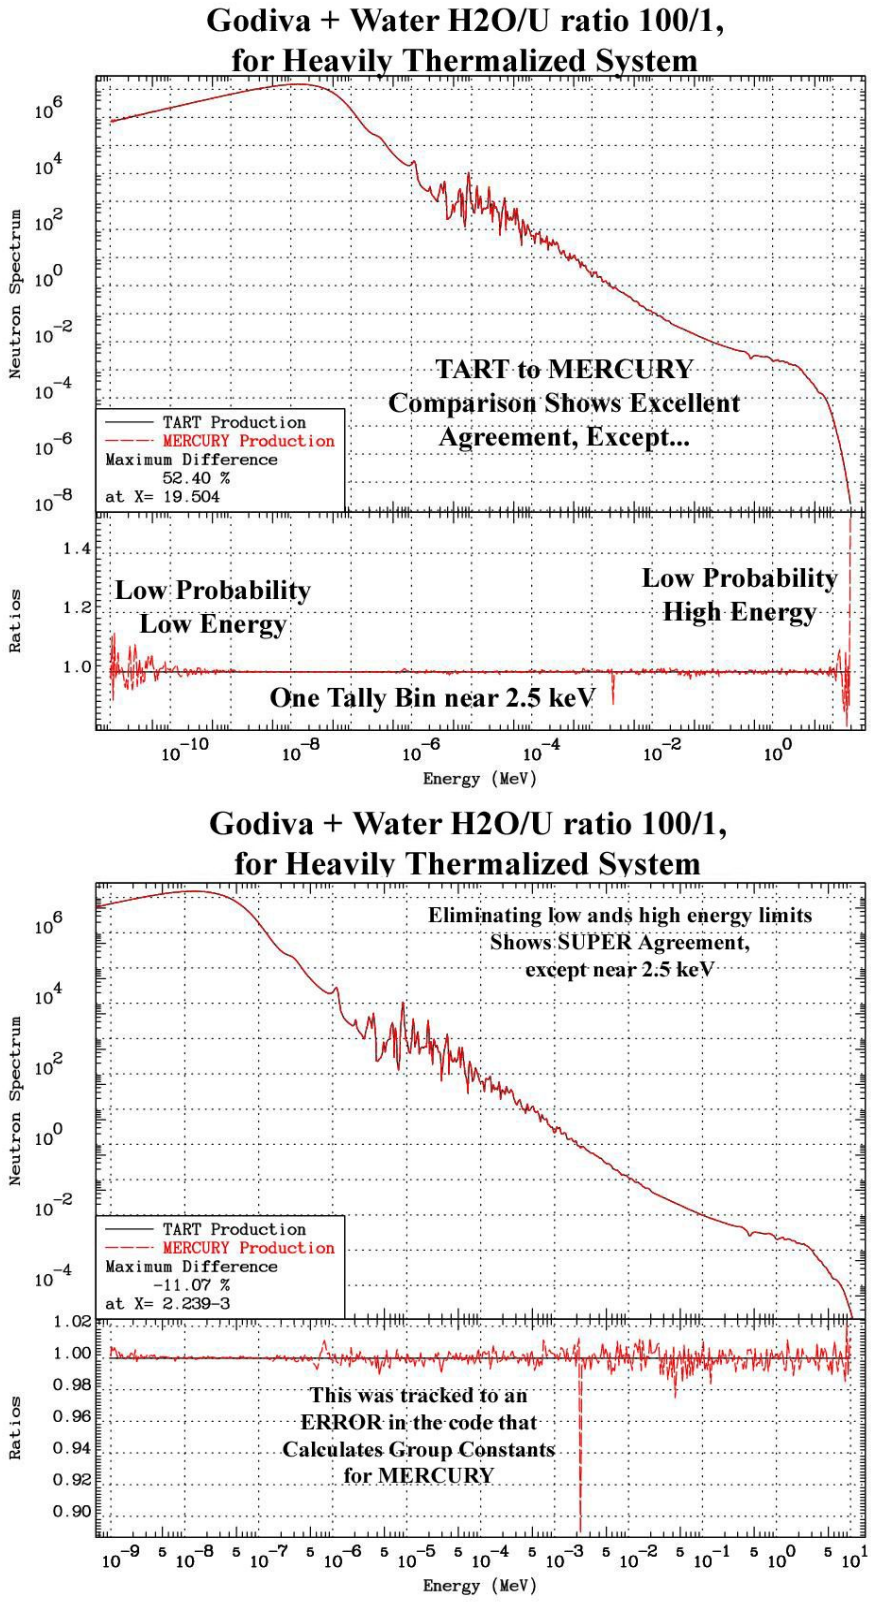

Figure 9. The particle production spectrum from the "Wet Godiva" criticality problem: (a) and (b) Excellent agreement between MERCURY (red) and TART (black) is observed, except in one energy bin near $E_n \simeq 1$ keV.

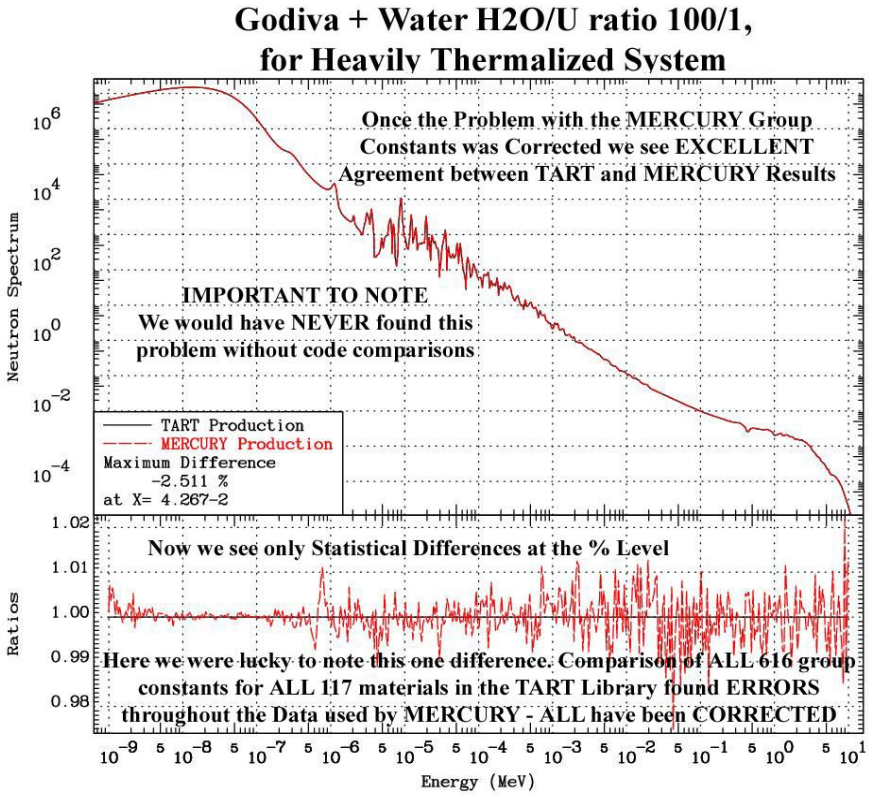

(**c**)

Figure 9 (continued). The particle production spectrum from the "Wet Godiva" criticality problem: (c) The modified version of MCFGEN which corrects a bug in the calculation of group constants for discontinuous point data yields excellent agreement between MERCURY (red) and TART (black).

The particle production spectrum from the Wet Godiva problem is shown in Figure 9. The MERCURY results shown in Figure 9 include all the modifications to the MCFGEN and MCAPM libraries which have been discussed above. The level of agreement between MERCURY (shown in red) and TART (shown in black) is excellent, except for statistical noise in the low-probability regions at either end of the energy range, and one group near $E_n \simeq 1$ keV.

This 10% variation in the spectrum is due to a 10% difference in the MERCURY group constant for ^{235}U in that one energy group, relative to the value used by TART. This difference in cross section values has been traced to the integration scheme which converts point wise cross section data to group constants in the MCFGEN. There is a discontinuity in the point wise data that falls within the boundaries of the energy group in question. The old, incorrect integration scheme used in MCFGEN produced a group constant that was 10% too large relative to the TART value. Once this integration scheme was recoded to account for discontinuities, the MERCURY results shown in Figure 9c were obtained.

It was fortuitous that this bug was uncovered. It was not observed in the fast version of Godiva (see Figure7), although the tell tale signs are there, buried in the statistical noise. When the entire set of MCAPM group constants were checked, another 828 examples of this effect were found in many of the isotopes in the data base. This bug suggests that one should have multiple copies of a given problem in the V&V test suite in order to probe the validity of the code in various energy regions (fast, intermediate, thermal, etc.).

6 SUMMARY AND FUTURE DIRECTIONS

The plan for, and initial results of, the Verification and Validation (V&V) of a modern, Monte Carlo particle transport code has been presented. In this paper, the MERCURY Monte Carlo code has been validated via code-to-code comparisons with TART, which itself has been previously validated against other Monte Carlo codes. This code intercomparison effort has focused on basic transport calculations using a Broomstick model, as well as criticality calculations of the k_{eff} eigenvalue for both the standard (fast) and "wet" (thermal) versions of the Godiva and Jezebel critical assemblies. Both integral results and particle spectra were used during these comparisons.

All new codes have their "growing pains". In the case of Mercury, the current effort has uncovered the following issues. The process of sampling of the fission (n, f) secondary neutron spectrum within the MCAPM library was found to be inaccurate due to (a) the lack of production of fission neutrons at intermediate to low energies ($E_n < 10$ keV), and (b) equally-probable sampling of the Watt spectrum at high energies ($E_n > 3$ MeV). Inaccurate sampling of other secondary particle spectra ($(n, n'), (n, 2n),$ etc.) was also found, which resulted from the lack of interpolation in the lowest-energy equally-probable bin. Finally, an incorrect integration scheme in the processing code MCFGEN that is used to convert the point wise, evaluated cross section and secondary distribution data into group constants, led to errors of 10% in certain group fluxes in MERCURY. All of these deficiencies and errors are corrected in the latest version of MERCURY.

Two key findings came out of the current work. The first, and perhaps most important conclusion, is that comparisons of integral parameters (such as $k_{eff}, \tau_{rem}, \alpha$, or energy deposits and dose rates) is necessary *but not sufficient* to guarantee that a transport code is working correctly. The second conclusion is that comparison of particle spectra is *invaluable*. Such comparisons uncovered *numerous* problems, and in many cases, was the *only* indication that there were underlying problems in the way that the nuclear data was being utilized.

This method of code intercomparison using detailed spectral information can be automated and completed within minutes of the transport calculation. In general, this method requires only a *small investment* of time and energy to implement the 616-group common tally structure. This effort cab results in a *big payoff* in terms of code reliability. At this writing, six codes are now using this common set of spectral tallies: MERCURY, TART, COG, VIM, KENO and MCNP (albeit, not in an automated fashion) now compare particle production, absorption, leakage and number density spectra. The authors encourage additional Monte Carlo codes to join this comparison network. You have everything to gain and nothing to lose!

In the near future, the V&V activities on the MERCURY code will be expanded to include several analytic benchmark problems in the test suite. These include subsets of (a) the Sood, Forster and Parson criticality problems, and (b) the Kobayashi, Sugimura and Nagaya deep-penetration transport problems. Several k_{eff} criticality benchmarks will also be added to the test suite, chosen from the ICSBEP Handbook, in order to test various fissile materials, geometries, and energy regimes.

The continuous-energy cross section capability in MERCURY will be validated via code comparison with a continuous energy Monte Carlo code, such as MCNP, as well as the multi-band, multi-group statistical treatment of resolved resonances which is used in TART. Finally,

the $S(\alpha,\beta)$ bound-atom (molecular) scattering model must be validated against the predictions of other codes for a variety of criticality problems. That will complete the validation of the prompt neutron transport capabilities within MERCURY.

7 ACKNOWLEDGMENTS

This work was performed under the auspices of the U.S. Department of Energy at the UC, Lawrence Livermore National Laboratory under Contract Number W-7405-Eng-48.

8 REFERENCES

1. D. Howe, "The Free On-line Dictionary of Computing", http://foldoc.doc.ic.ac.uk/foldoc/Dictionary.gz (2004).
2. R. J. Procassini and J. M. Taylor, *MERCURY User Guide (Version b.6)*, Lawrence Livermore National Laboratory, Report UCRL-TM-204296 (2004).
3. D. E. Cullen, *TART 2002: A Coupled Neutron-Photon 3-D, Combinatorial Geometry Time Dependent Monte Carlo Transport Code*, Lawrence Livermore National Laboratory, Report UCRL-ID-126455, Revision 4 (2002).
4. R. Buck, E. Lent, T. Wilcox and S. Hadjimarkos, *COG User's Manual (Fifth Edition) A Multiparticle Monte Carlo Transport Code*, Lawrence Livermore National Laboratory, Internal Report (2002).
5. A. Sood, R. A. Forster and D. K. Parsons, *Analytic Benchmark Test Set for Criticality Codes*, Los Alamos National Laboratory, Report LA-UR-01-3082 (2001).
6. K. Kobayashi, N. Sugimura and Y. Nagaya, *3-D Radiation Transport Benchmark Problems and Results for Simple Geometries with Void Regions*, Nuclear Energy Agency, Report NSC-DOC2000-4 (2000).
7. International Criticality Safety Benchmark Evaluation Program, *International Handbook of Evaluated Criticality Safety Benchmark Experiment*, Nuclear Energy Agency, Report NEA/NSC/DOC(95)03/I (2003).
8. L. F. Hansen, T. Komoto, E. F. Plechaty, B. A. Pohl, G. S. Sidhu and C. Wong, "Pulse Spheres", *Nucl. Sci. and Eng.*, **62**, p. 550 (1977).
9. L. J. Cox, et al., *MCNP Version 5*, Los Alamos National Laboratory, Report LA-UR-02-1527 (2002).
10. D. E. Cullen, R. N. Blomquist, C. Dean, et al., *How Accurately can we Calculate Thermal Systems?*, Lawrence Livermore National Laboratory, Report UCRL-TR-203892 (2004).
11. P. B. Fox and L. M. Petri, *Validation and Comparison of KENO V.a and KENO-VI*, Oak Ridge National Laboratory, Report ORNL/TM-2001/110 (2002).
12. R. N. Blomquist, *Status of the VIM Monte Carlo Neutron/Photon Transport Code*, Argonne National Laboratory, Report ANL/RAE/CP-106063 (2002).
13. D. E. Cullen, *PREPRO 2002: 2002 ENDF/B Pre-processing Codes*, International Atomic Energy Agency, Report IAEA-NDS-39, Revision 11 (2002).
14. P. S. Brantley, C. A. Hagmann and J. A. Rathkopf, *MCAPM User Guide (Version 1.1)*, Lawrence Livermore National Laboratory, Internal Report (2002).